

EMBEDDED REVOCATION MESSAGING

This application relates to the information processing arts. It finds particular application in the control of dissemination and use of copyrighted music and other recorded content through the use of device revocation lists, and will be described with particular reference thereto. However, the following will also find application more generally in
5 controlling information transfer among networked devices, in providing improved property protection of music, video, software, data, and other information content, in combating proliferation of viruses and other malicious software code, and in other similar applications.

Revocation lists are known for controlling the spread of viruses, for limiting the distribution and use of copyright protection-disabling software, excluding devices that have
10 been tampered with, and the like. Security protocols that employ revocation lists include the X.509 authentication structure and the DTCP (also known as 5C) IEEE1394 link security protocol. The revocation list identifies networked machines which are regarded as security risks or are otherwise undesirable. Networked devices receive the revocation list and henceforth refuse to communicate with devices identified on the revocation list. Preferably, the
15 revocation list is updated or supplemented on an occasional basis to account for spreading viruses, newly identified problematic devices, or the like.

Revocation lists provide relatively straightforward and effective protection against computer viruses and other information content which is widely regarded as undesirable. Revocation lists are also employed to exclude a user from having access to certain
20 network resources, or to exclude devices from joining a communication. In the computer environment, the users of networked devices generally are cooperative in receiving and storing the revocation lists and the occasional updates or supplements to the revocation lists.

However, in certain areas, especially in the areas of music and video copyright protection and other property protection schemes, a problem arises in using revocation lists to
25 enforce such property rights. Certain users of networked devices who want to circumvent copyright or other property protection are unlikely to cooperate in receiving the revocation data. Rather, these users are likely to attempt to block distribution of revocation lists, or to attempt to delete revocation lists that are stored on networked devices. Moreover, even users who are indifferent or favorably disposed toward the revocation lists may not receive or retain
30 them due to the activities of viruses or other malicious software which act to block

transmission and/or storage of the revocation lists.

In the past, hostility and active opposition to revocation list distribution has been countered using various approaches. In one typical revocation list distribution method, when a connection between two devices is established, the network connection protocol calls
5 for and requires communication and storage of the revocation list (if a new or updated revocation list is being distributed) before any other action is allowed. In this way, devices are forced to receive a revocation list. A device having received a revocation list will store it and subsequently refuse to communicate with devices on that revocation list.

However, such secured network transmission protocols are not immune to the
10 countermeasures of hackers who continually work toward developing new techniques for overcoming these network protocol-based revocation list distribution techniques. Efforts of hackers may be successful when they can readily identify the revocation list in the data stream, and because the hackers (or users of devices which have been infected by a virus or other malicious program produced and distributed by a hacker) experience no adverse consequences
15 when the revocation list is removed from the data stream.

The present invention contemplates an improved apparatus and method that overcomes the aforementioned limitations and others.

|

20 According to one aspect, an apparatus is disclosed, including a means for storing, transmitting, or receiving a signal representing user-desired content. The signal includes the user-desired content, and device revocation information embedded in the user-desired content.

25 According to another aspect, a method is provided for distributing revocation information. The revocation information is embedded into user-desired content.

One advantage resides in integrating a revocation list with content desired by the user. In this manner, the user is induced to accept the revocation list as part of the desired content.

Another advantage resides in its capacity for high rate systems to communicate
30 large revocation lists which can contain identification information pertaining to a large number of devices.

Yet another advantage resides in distributing the revocation message over a large portion of musical content or other user-desired content. In this way, the user cannot avoid receiving the revocation message by omitting a small part of the musical or other user-desired content.

5 Numerous additional advantages and benefits will become apparent to those of ordinary skill in the art upon reading the following detailed description of the preferred embodiments.

10 The invention may take form in various components and arrangements of components, in software, and in various process operations and arrangements of process operations. The drawings are only for the purpose of illustrating preferred embodiments and are not to be construed as limiting the invention.

FIGURE 1 shows an exemplary three networked devices and schematically
15 illustrates transfer of a revocation list between two devices, and subsequent refusal of content transfer to the third device which is included in the revocation list.

FIGURE 2 schematically shows components of the device 1 that encode the user-desired content along with an embedded revocation list.

FIGURE 3 schematically shows components of the device 2 that recover the
20 revocation list and optionally decode the user-desired content.

FIGURE 4 schematically shows alternative components of the device 1' that encode user-desired content along with an embedded revocation list.

FIGURE 5 schematically shows an optical compact disk which includes a digital watermark with a revocation list.

25 FIGURE 6 shows an exemplary method for manufacturing the optical compact disk of FIGURE 5.

30 With reference to FIGURE 1, devices 1, 2, 3 communicate audio, video, or other content with one another by a suitable transfer protocol. Such transmission can be by an electronic or optical digital communication network, a wireless network such as Bluetooth, or

the like. The device 1 has recently received an updated revocation list 4 that includes a revocation of the device 3.

The revocation list 4 can specify the device 3 in a variety of ways. The device 3 can be specifically revoked, that is, the device 3 is uniquely identified in the revocation list 4 as a revoked device. Alternatively, a model or manufacturer corresponding to the device 3 can be revoked. In yet another revocation arrangement, a content transfer protocol, pathway, or the like used to connect the device 3 with the devices 1, 2 can be revoked.

The device 1 communicates content 5 to the device 2. The communicated content 5 includes the revocation list 4 embedded in the content 5 in such a way that the embedded revocation list 4 is not easily removed from the content 5 or is removable only with substantial degradation of the content. Moreover, the revocation list 4 is preferably divided into sub-lists which are repeated throughout the length of content 5 so that the revocation list 4 is substantially received at the device 2 even if a small portion of the content 5 is discarded.

Subsequent to communication of the content 5, an attempt is made to initiate a transfer of content 6 from the device 2 to the device 3. However, the device 2 does not transfer content 6 in view of the updated revocation list 4 which the device 2 received in embedded form from the device 1 during transfer of content 5. The revocation list 4 identifies the device 3 as a revoked device (or identifies the device 3 as communicating with a revoked communication protocol or pathway, etc.), and so the device 2 does not communicate with the device 3.

Although audio content 5, 6 is described, the content can also be video content, data content, software, or the like. Moreover, the information can be represented or encoded in any suitable format, such as PCM, MPEG, AC3, DST, MLC, ATRAC, DivX, analog, or the like. The revocation list 4 can be embedded in the content in a suitable manner, such as by an embedded watermark, a designated audio or video channel, in a physical watermark, embedding in an analog signal, and so forth. The choice of revocation list embedding technique will depend on the type of content (audio, video, software, etc.), and the content format.

With reference to FIGURE 2, in one exemplary approach the revocation list is embedded in DST encoded audio content 5 via a digital watermark that contains the revocation list 4. The watermark optionally contains additional content such as copyright information. The digital watermark does not affect the audio or other user-desired content, but also is not

readily removed from the content stream 5 without decoding and reencoding the stream. In this exemplary embodiment, the device 1 performs lossless DST encoding of a high fidelity audio stream 10. Encoding of audio is shown by way of illustrative example, although encoding of multimedia, video, data, and the like are also contemplated. During the encoding, the revocation list 4 is embedded. The audio stream can include a plurality of audio channels, for example to transmit stereo or surround-sound audio. Moreover, instead of or in addition to an audio stream, the user-desired content could include video content, data content, or the like.

A framing processor 14 frames the audio stream 10 into data frames for lossless encoding. Preferably, if the audio stream 10 includes a plurality of audio channels, each audio channel is separately framed by the framing processor 14. In one suitable embodiment, a constant frame length of 37,362 bits per frame is used. However, different frame lengths can be used to optimally balance encoding performance and other factors. For each frame output by the framing processor 14, an encoding parameters processor 16 computes suitable encoding parameters. In typical lossless encoding schemes, the encoding parameters processor 16 computes predictive parameters that are used in the encoding to approximate the frame contents. Various filtering processes are known in the art for generating a good set of predictive parameters. Such filtering is computationally intensive, and the encoding performance is not critically dependent upon using fully optimized predictive parameters. Hence, typically the encoding parameters processor 16 computes predictive parameter values that provide efficient, but not optimal encoding performance.

Because precisely optimized values of the predictive parameters is not critical, these parameters can be modified to encode watermark content 20 without substantially degrading the efficiency of the lossless encoding. The watermark content 20 includes at least the revocation list 4. Preferably, a sub-lists generator 18 divides the revocation list into sub-lists that are distributed through the watermark content 20, and hence are ultimately distributed through the DST content stream 5. The sub-lists generator 18 optionally also duplicates the revocation list (or sub-lists thereof) to ensure that the content stream 5 has embedded revocation information extending throughout the content stream 5.

A watermark encoder 22 modifies the predictive parameters in a predetermined manner to encode the watermark content 20. For example, a least significant bit of one or more of the predictive parameters can be set to one or zero corresponding to one bit of the watermark content 20. This small change in one or a few parameter values generally does not

significantly change the subsequent lossless encoding efficiency. In another suitable modification, an additional dummy predictive parameter is selectively added so that the number of predictive parameters (even or odd) corresponds with one bit of the watermark content 20. Similarly, in yet another embodiment one predictive parameter is selectively
5 deleted so that the number of predictive parameters (even or odd) corresponds with one bit of the watermark content 20. Other predetermined modifications of the predictive parameters can be used to encode the watermark content 20.

A lossless frame encoder 24 encodes the frame using the modified predictive parameters output by the watermark encoder 22. In one suitable lossless encoding scheme, a
10 residual is computed that corresponds to a difference between the frame value and a value predicted using the predictive parameters. For good predictive parameters, this residual contains much less relevant information, and can be efficiently encoded. The residual is suitably entropy-encoded. The encoded content along with the predictive parameters are arranged into a lossless coded frame by a lossless coded frame constructor 26. Preferably, the
15 predictive parameters are processed by a compressor 28 using an efficient compression algorithm prior to incorporation into the lossless coded frame. The lossless coded frame includes the compressed predictive parameters, the encoded content output by the frame encoder 24, and suitable control information (for example a total number of predictive parameters and/or a frame length) in a predetermined format. The lossless coded frame enters
20 the DST stream 5 for transmission to the device 2.

With reference to FIGURE 3, the DST stream 5 is received at the device 2 and processed by a lossless coding frame processor 30 that extracts the lossless coding frames from the DST content stream 5 for decoding. A frame components analyzer 32 separates out the compressed predictive parameters and the lossless encoded data components. The
25 compressed predictive parameters are transmitted to an encoding parameters decompressor 34 for decompression to recover the predictive parameter values. A watermark content extractor 36 recovers the watermark content 20 including the revocation list 4 from the predictive parameters. The watermark content 20 is recoverable without decoding the audio content. This allows the revocation list to be recovered even if the content stream 5 is stored in encoded
30 form at the device 2 without immediately decoding and playing the audio content. Moreover, if the transfer rate of the content stream 5 exceeds the decoding speed, the watermark content can be extracted while the audio content is buffered for decoding.

To recover the audio content, a lossless frame decoder 40 decodes the encoded data components using the predictive parameters recovered by the parameters decompressor 34 to recover the audio stream 10.

The digital watermark content 20 including the revocation list 4 are logically
5 integrated with the high fidelity audio stream 10. Although the DST digital watermark is typically regarded as a fragile watermark, those skilled in the art will recognize that in practice modifying the integrated DST data stream to remove the watermark while retaining a valid high fidelity audio DST stream involves a substantial amount of time, effort, and cost. Moreover, the revocation list 4 is preferably broken into independently readable sub-lists, each
10 of which is integrated into one or a few lossless coded frames. In this way, the revocation list cannot be avoided by omitting a small amount of the content stream 5.

With reference to FIGURE 4, another suitable device 1' embeds watermark content 20 including the revocation list 4 with a multi-channel audio stream. The audio stream includes audio channel 0 10', audio channel 1 10'', and so forth through audio channel n 10'''.
15 Each audio channel 10', 10'', 10''' is processed by a framing processor 14', 14'', 14''' and lossless encoded through the entropy encoding stage by lossless encoders 44', 44'', 44'''. As is known in the art, entropy encoding typically outputs the residual (e_i) and a probability signal (p_i). The residual and probability signals are multiplexed and arithmetically encoded by a MUX/arithmetic encoder 46. The predictive parameters are compressed by a compressor 28',
20 and combined with the multiplexed and encoded data into a lossless coded frame by a lossless coded frame constructor 48 for transfer via the content stream 5'.

To embed the watermark 20 including the revocation list 4, at least one binary bit of watermark content 20 is included as an input to the MUX/encoder 46. For encoders that require paired bit input data, the watermark bit is preferably paired with one of the probability
25 signal bits, for example paired with the bit p_n . At the receiving device, the DST stream 5' is demultiplexed and arithmetically decoded to recover the watermark 20. As with the approach of FIGURES 2 and 3, the content stream 5' does not need to be lossless decoded to recover the revocation list 4.

DST digital watermarks provide a relatively large messaging bandwidth. For an
30 exemplary 75 frames/second transfer rate with 15 bits encoded per frame (for example, corresponding to an audio stream 10 including fifteen lossless encoded audio channels with

each channel having one embedded watermark bit) a watermark transfer data rate of about 1 kbit/sec is obtained.

In a three minute song, this results in a corresponding digital watermark capacity of 180 seconds times 1 kbit/second = 180 kbits or about 22 kbytes. For a typical device identifier length of 8 bytes and a signature length of 128 bytes (that is, a signature of a license authority which is typically required to demonstrate authenticity of the distributed revocation list) the portion of the digital watermark corresponding to a three minute song can encode a revocation list identifying approximately 2800 devices (that is, 22 kbytes divided by 8 bytes per device identifier).

Embodiments including embedding of revocation list information in a watermark of a DST encoded audio stream, and embedding of revocation list information using a designated audio channel of an encoded multi-channel audio stream are described herein by way of example. However, those skilled in the art can readily adapt the described embodiments for use with other types of information content and other types of encoding. For example, the embedded watermark approach is readily adapted for revocation list distribution in watermarks of encoded video streams. Similarly, an unused audio or video channel of a multimedia content stream can be used to transmit embedded revocation list information. Software and data compression technologies can similarly incorporate embedded revocation lists. Revocation list information can be incorporated as a dedicated "mute" audio channel of DVD video. The mute channel is encoded using PCM, MPEG, AC3, or another audio encoding supported by DVD, and is ignored when the DVD is played. However, this approach has a disadvantage if the encoding is computationally complex in that the watermark is usually not recoverable without fully decoding the audio content. Meridian lossless packing (MLP) lossless encoding, which has been selected for DVD-audio, is readily adapted for embedding a revocation list.

In these embedding methods, the revocation list is preferably divided into revocation device sub-lists that are distributed substantially co-extensively with the video or audio content. The lists can be redundantly repeated multiple times so that the user cannot avoid the revocation list by omitting receipt, playback or other accessing of a limited portion of the user-desired programming.

With reference to FIGURE 5, yet another exemplary embodiment, in which a revocation list is embedded in a digital watermark of an optical compact disk 60, is described.

The optical compact disk 60 can be a DVD type of disc including high fidelity audio content 62 that is encoded by optically reflective pits disposed on or in the disk 60. The reflective pits encoding the high fidelity content 62 extend across and substantially fill a central portion of the compact disk 60. The region of high fidelity audio content 62 is represented by texturing in
5 FIGURE 5. In a suitable encoding, the high fidelity audio content 62 is encoded in the lengths of the optically reflective pits while the digital watermark 64 is encoded in pit widths.

The watermark 64 is preferably substantially coextensive with the high fidelity audio content 62. In FIGURE 5, an exemplary spatial extent of encoded data of the digital watermark 64 on the disk 60 is schematically indicated by inner and outer dashed
10 circumferential boundary lines. The watermark 64 includes a revocation list which is preferably divided into redundantly replicated sub-lists distributed throughout the disk 60. The watermark 64 optionally also includes other information such as copyright information.

In operation, a user plays high fidelity audio content 62 of the compact disk 60 on a playback device (not shown). If the playback device is a compliant device respective to
15 the revocation list, the playback device simultaneously reads the watermark 64 during playback of the high fidelity audio content 62. If the revocation list contained in the watermark 64 post-dates any revocation lists currently stored on the compliant playback device, the revocation list of the watermark 64 replaces the older stored revocation list.

Although an exemplary optical disk 60 is described here, revocation
20 information can similarly be distributed by distributing other non-volatile storage media on which is stored user-desired content with the embedded revocation information. For example, solid-state non-volatile memory units, magnetic disks, and the like on which the user-desired content with embedded revocation information is stored can be distributed.

With reference to FIGURE 6, a suitable method 70 for mass-producing the
25 optical compact disk 60 is described. A revocation list 71 is divided 72 into a plurality of revoked device sub-lists 74. The device sub-lists 74 are optionally combined with copyright information 75 or other additional watermark content, and the watermark information is encoded 76 and a target reflective pit width is computed 78. Similarly, the high fidelity audio content 62 is encoded 80 and a target reflective pit length is computed 82.

30 The computed target pit length and width are used to modulate 86 a laser that records the reflective pits that encode the high fidelity audio data stream. The pits are formed 88 on a glass substrate disk to define a glass master 90. The glass master 90 is used to

mass-produce 92 commercial optical compact disks, such as the disk 60, which are usually plastic disks with reflective optical coatings. Optionally, the mass-production 92 further includes formation of a conventional redbook layer for backward compatibility with older CD players.

5 Preferably, the revocation list 71 is updated on an occasional or more frequent basis by producing an updated glass master 90 in accordance with the relevant process operations of the method 70. The revocation list 71 is updated, for example, to include new devices that have been identified as being non-compliant with the selected content reproduction management protocol. In this manner, the compact disk manufacturer continually
10 produces compact disks with music offerings that contain up-to-date revocation lists. Compliant devices automatically receive an updated revocation list each time a new compact disk containing updated or supplementary revocation list information is played.

Embodiments embedding revocation information into audio streams and into physical medium watermarks have been described. However, it will be appreciated that
15 revocation information can be similarly embedded and distributed by other devices or media which transmit, receive, or store various types of user-desired audio, video, data, software, or other content. Employment of less fragile digital watermarks than the DST watermarks described herein makes it more difficult to remove the embedded revocation list. However, less fragile digital watermarks typically have less data capacity and hence can store fewer
20 revoked device identifications. Moreover, instead of using a digital watermark, the revocation information can be encoded in a mute channel or other pathway provided by the encoding technology and comporting with the type of content being processed.

The invention has been described with reference to the preferred
embodiments. Obviously, modifications and alterations will occur to others upon reading
25 and understanding the preceding detailed description. It is intended that the invention be construed as including all such modifications and alterations insofar as they come within the scope of the appended claims or the equivalents thereof.